Accelerating Science on Titan: Highlights of new HEP and NP Projects

20 Years of Excellence in Computational Science
OLCF
OAK RIDGE LEADERSHIP COMPUTING FACILITY
1992–2012

Kenneth Read
Oak Ridge National Laboratory/
University of Tennessee

BNL HPC 2013
June 5, 2013

What is the Leadership Computing Facility?

- Collaborative DOE Office of Science program at Oak Ridge and Argonne National Laboratories
- Mission: Provide the computational and data science resources required to solve the most important scientific & engineering problems in the world.
- Highly competitive user allocation programs (INCITE, ALCC).
- Projects receive 10x to 100x more resource than at other generally available centers.
- LCF centers partner with users to enable science & engineering breakthroughs.

Argonne NATIONAL LABORATORY

OAK RIDGE National Laboratory

Kepler GK110 GPU

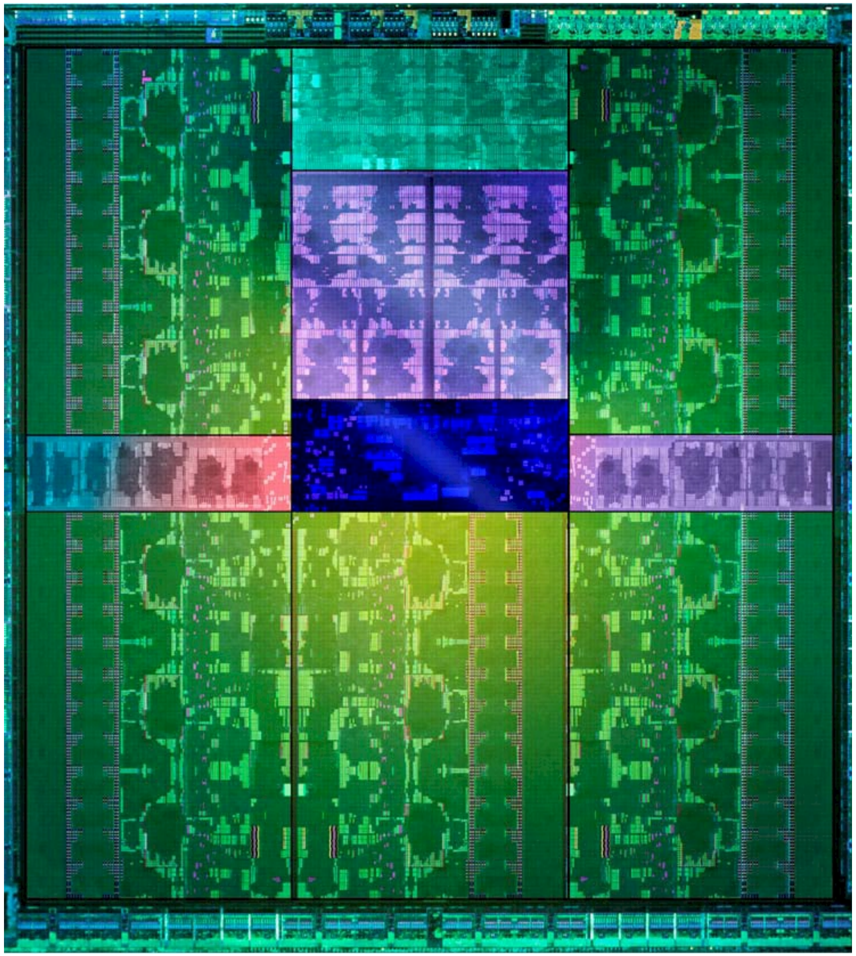

2.3 cm

Most complex semiconductor device ever.
Delivers 1.3 TFlop peak double precision.

NVIDIA Tesla Kepler K20X

NVIDIA GeForce GTX Titan – On Sale Now

Science breakthroughs at the LCF:

A few of the many science and engineering advances through the INCITE program

Hours requested vs. allocated: ***~2X per year*** (2007–2010) ***~3X per year*** (2011–2013)

	2004	2005	2006	2007	2008	2009	2010	2011	2012	2013
Hours allocated	4.9M	6.5M	18.2M	95M	268M	889M	1.6B	1.7B	1.7B	***5B***
Projects	3	3	15	45	55	66	69	57	60	61

Researchers solved the 2D Hubbard model and presented evidence that it predicts HTSC behavior, ***Phys. Rev. Lett*** *(2005)* .

Modeling of molecular basis of Parkinson's disease named #1 computational accomplishment, ***Breakthroughs*** *(2008)*.

Largest simulation of a galaxy's worth of dark matter, showed for the first time the fractal-like appearance of dark matter substructures, ***Nature*** (2008), ***Science*** *(2009)*.

World's first continuous simulation of 21,000 years of Earth's climate history, ***Science*** *(2009)*.

Largest-ever LES of a full-sized commercial combustion chamber used in an existing helicopter turbine, **Compte Rendus de Mecanique** (2009).

Unprecedented simulation of magnitude-8 earthquake over 125-square miles, ***Proceedings SC10.***

NIST proposes new standard reference materials from LCF concrete simulations, ***Eur Phys J E Soft Matter*** *(2012)*.

Calculation of the number of bound nuclei in nature, ***Nature*** *(2012)*.

New method to rapidly determine protein structure, with limited experimental data, ***Science*** *(2010)*, ***Nature*** *(2011)*.

OMEN breaks the petascale barrier using more than 220,000 cores, ***Proceedings SC10.***

Science requires exascale capability in this decade

Mission: Deploy and operate the computational resources required to tackle global challenges

- Deliver transforming discoveries in climate, materials, biology, energy technologies, etc.
- Enabling investigation of otherwise inaccessible systems, from regional climate impacts to energy grid dynamics

Vision: Maximize scientific productivity and progress on largest scale computational problems

- World-class computational resources and specialized services for the most computationally intensive problems
- Stable hardware/software path of increasing scale to maximize productive applications development

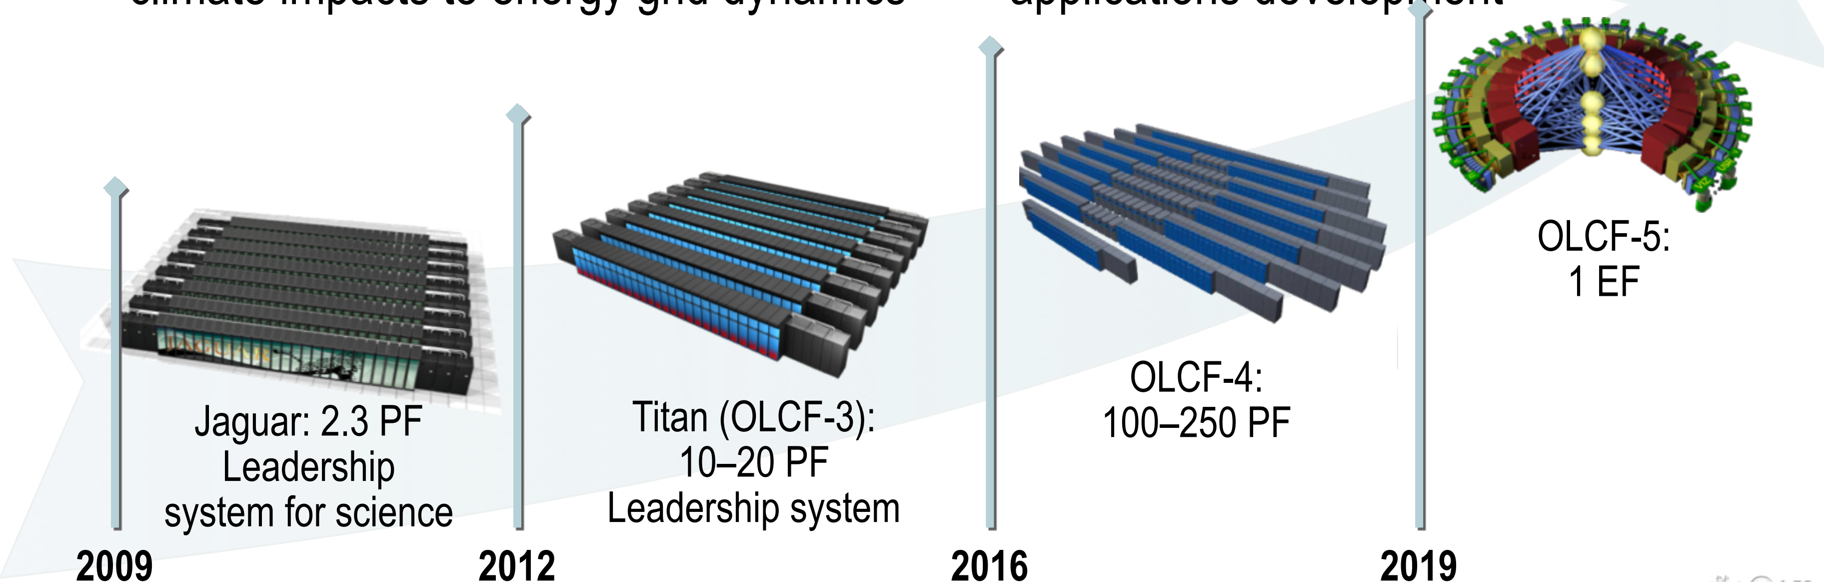

Slide courtesy of Jeff Nichols

Architectural Trends – No more free lunch

- Moore's Law continues (green) but CPU clock rates stopped increasing in 2003 (dark blue) due to power constraints (blue).
- Power is capped by heat dissipation and $$$. Confronting the *power wall*.
- Performance increases have been coming through increased parallelism.

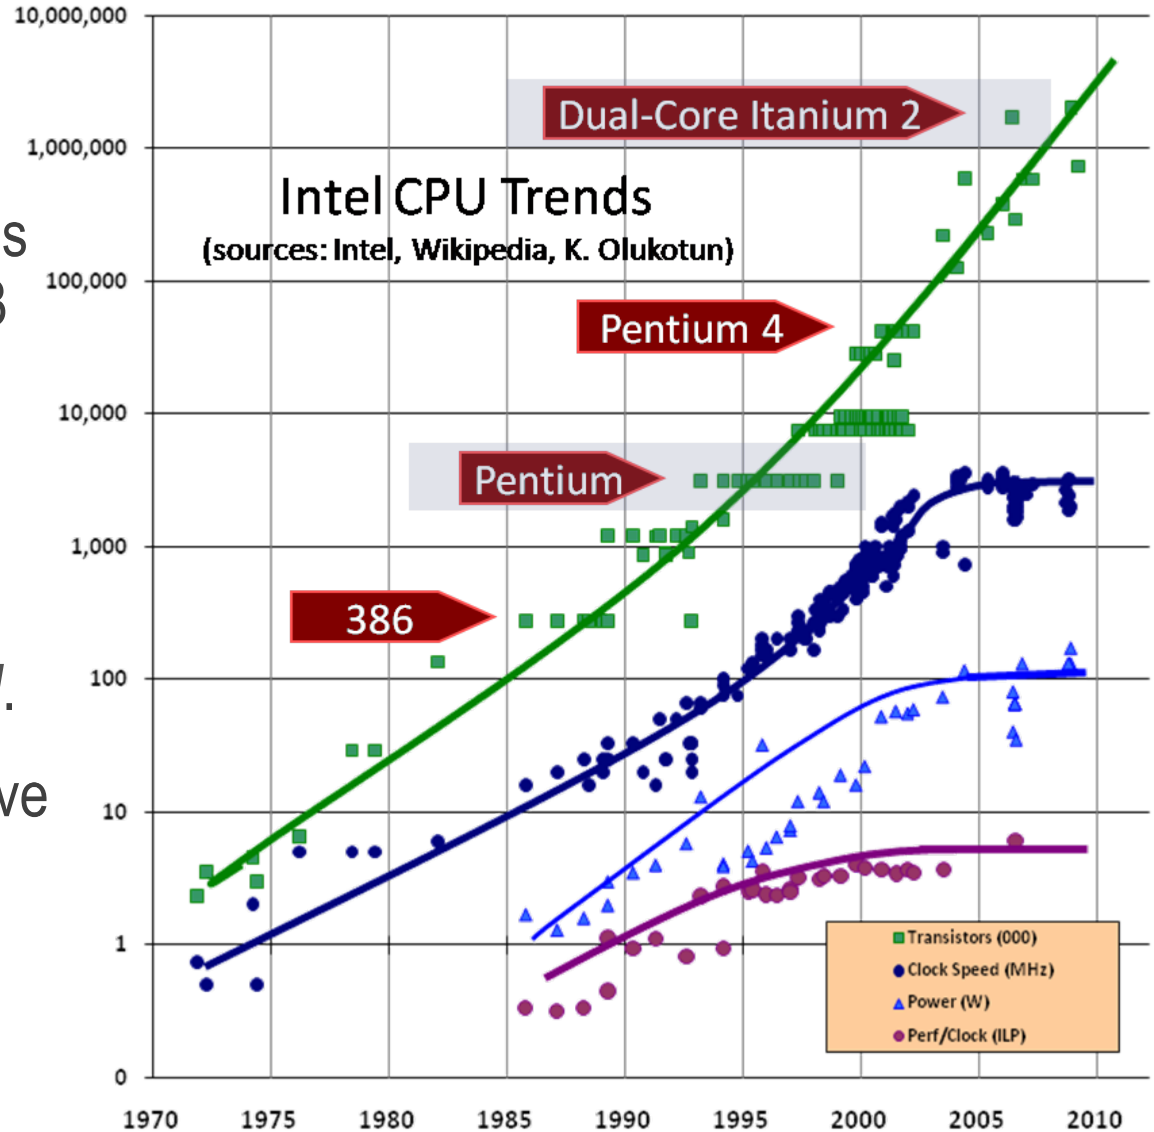

Herb Sutter, Dr. Dobb's Journal:
http://www.gotw.ca/publications/concurrency-ddj.htm

Power is THE problem

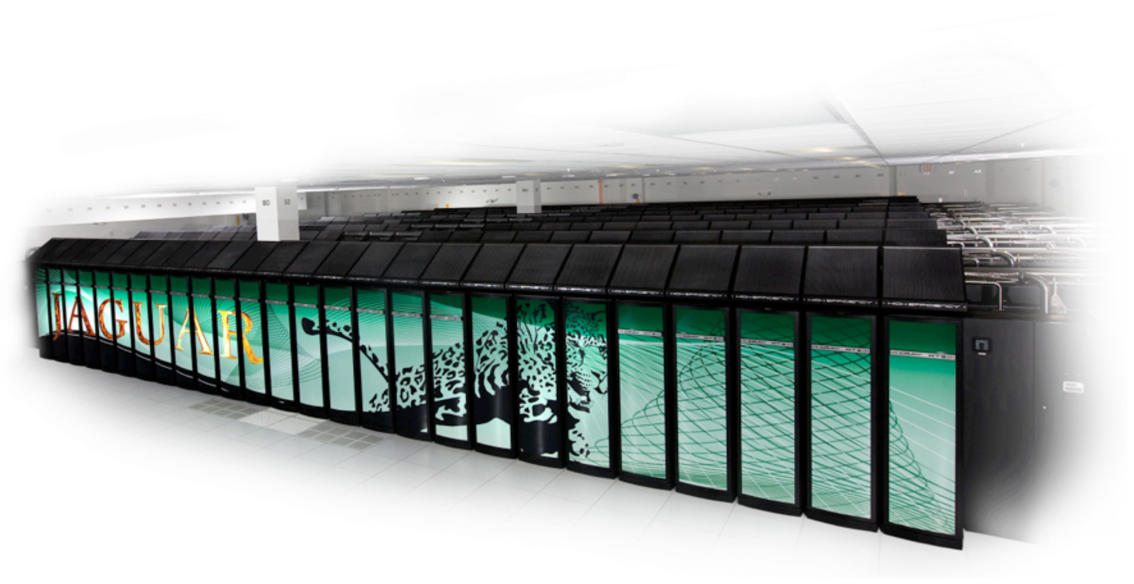

Power consumption of 2.3 PF Jaguar:
7 megawatts, equivalent to that of a small city (5,000 homes)

Scaling via traditional CPUs is no longer economically feasible

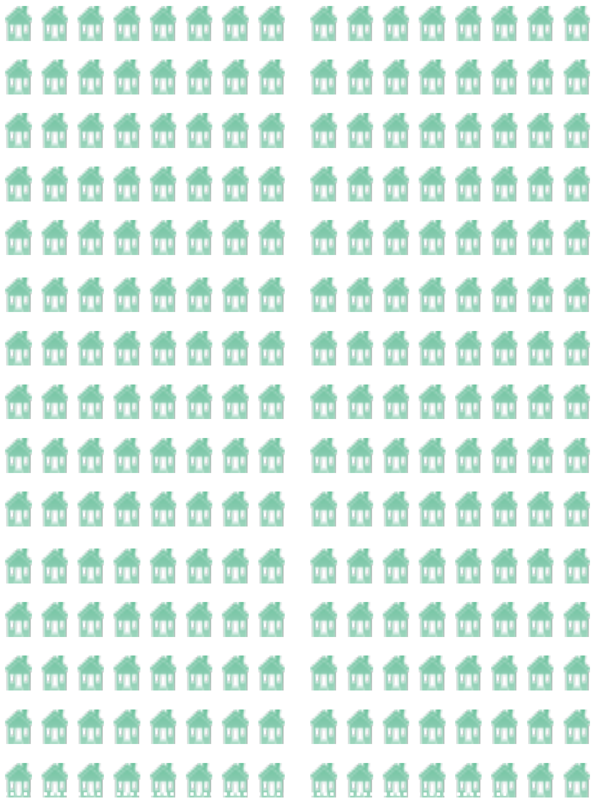

20 PF+ system:
30 megawatts (30,000 homes)

Why GPUs?
High performance & power efficiency towards exascale

- Hierarchical parallelism improves scalability of applications
- Expose more parallelism through code refactoring and source code directives
  - Doubles performance of many codes
- Heterogeneous multicore processor architecture: Using right type of processor for each task
- Data locality: Keep data near processing
  - GPU has high bandwidth to local memory for rapid access
  - GPU has large internal cache
- Explicit data management: Explicitly manage data movement between CPU and GPU memories

CPU

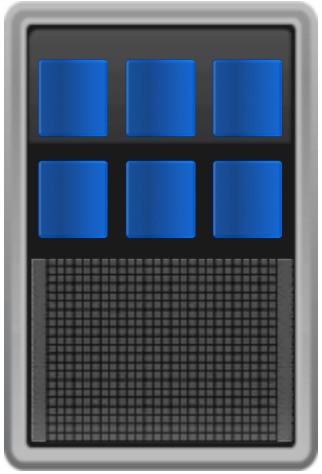

- Optimized for sequential multitasking

GPU Accelerator

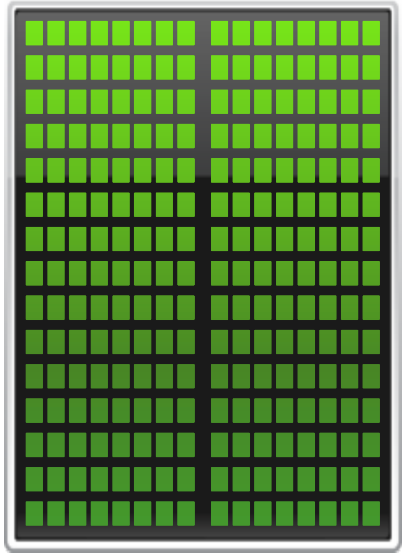

- Optimized for many simultaneous tasks
- 10× performance per socket
- 5× more energy-efficient systems

	Titan Nodes (Cray XK7)		
Node	AMD Opteron 6200 Interlagos (16 cores)	2.2 GHz	32 GB (DDR3)
Accelerator	Tesla K20x (2688 CUDA cores)	732 MHz	6 GB (DDR5)

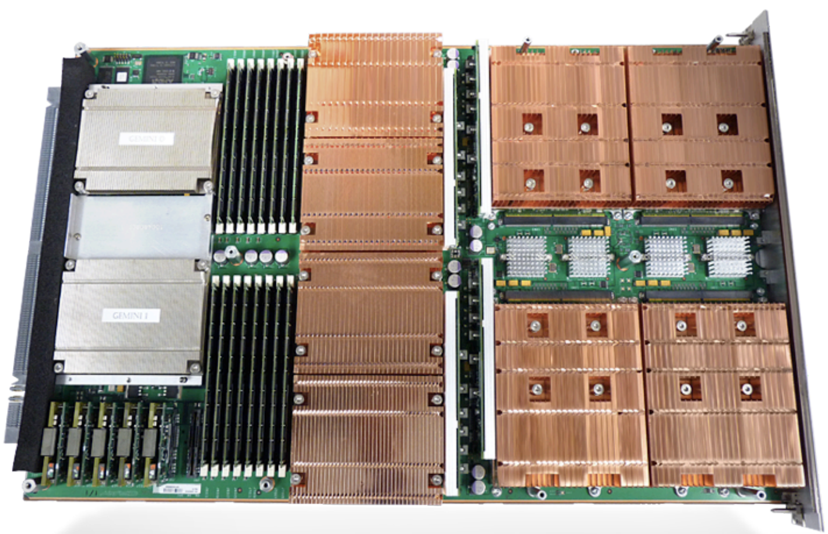

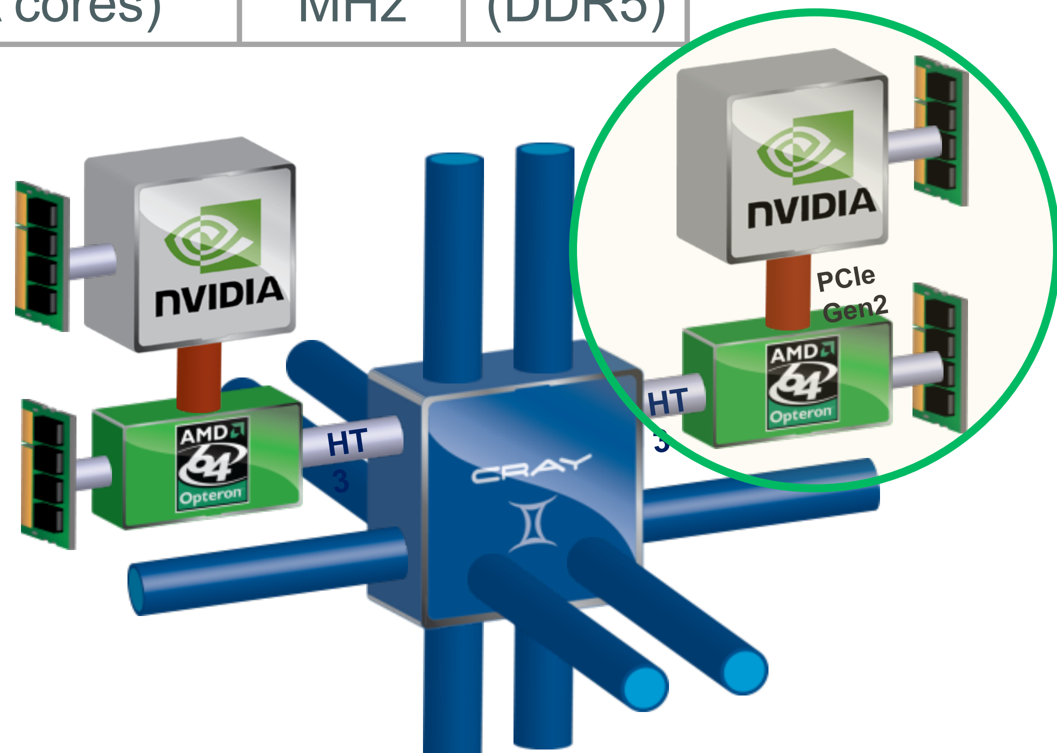

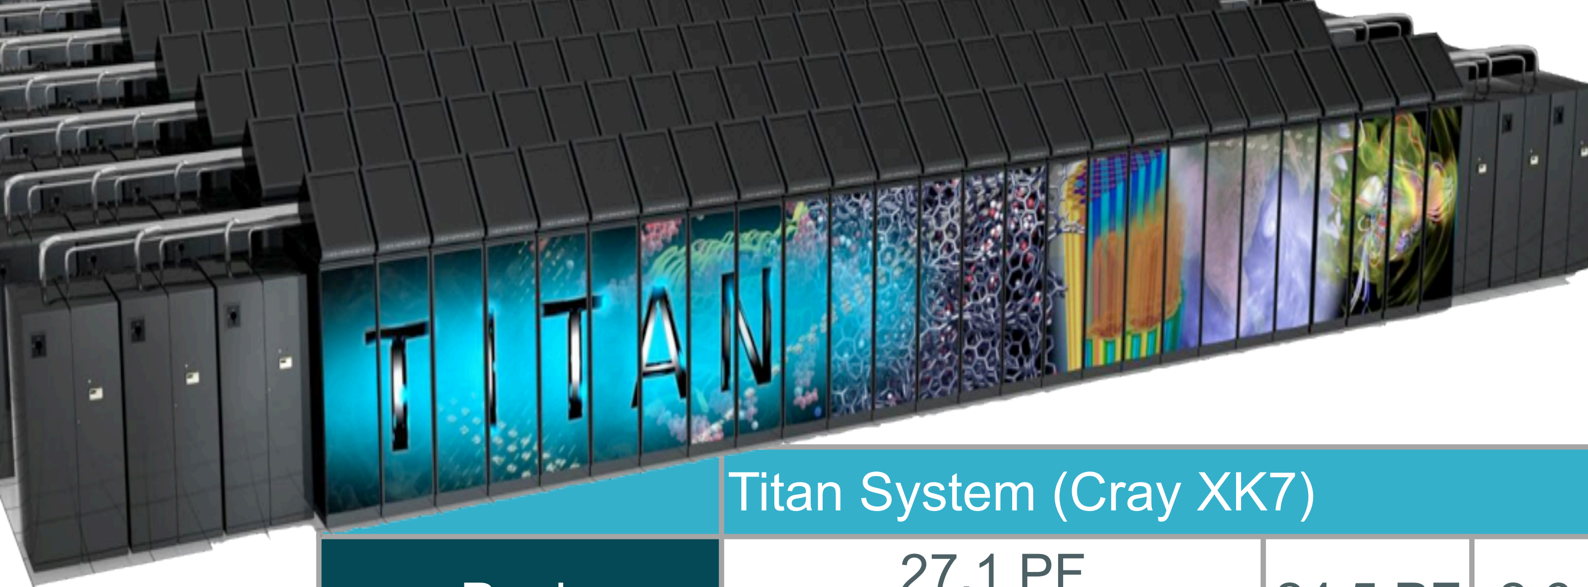

	Titan System (Cray XK7)		
Peak Performance	27.1 PF 18,688 compute nodes	24.5 PF GPU	2.6 PF CPU
System memory	710 TB total memory		
Interconnect	Gemini High Speed Interconnect	3D Torus	
Storage	Lustre Filesystem	32 PB	
Archive	High-Performance Storage System (HPSS)	29 PB	
I/O Nodes	512 Service and I/O nodes		

#1 TOP500 SUPERCOMPUTER SITES

17.59 PF
8.2 Megawatts

THE GREEN 500 #3

Hybrid Programming Model

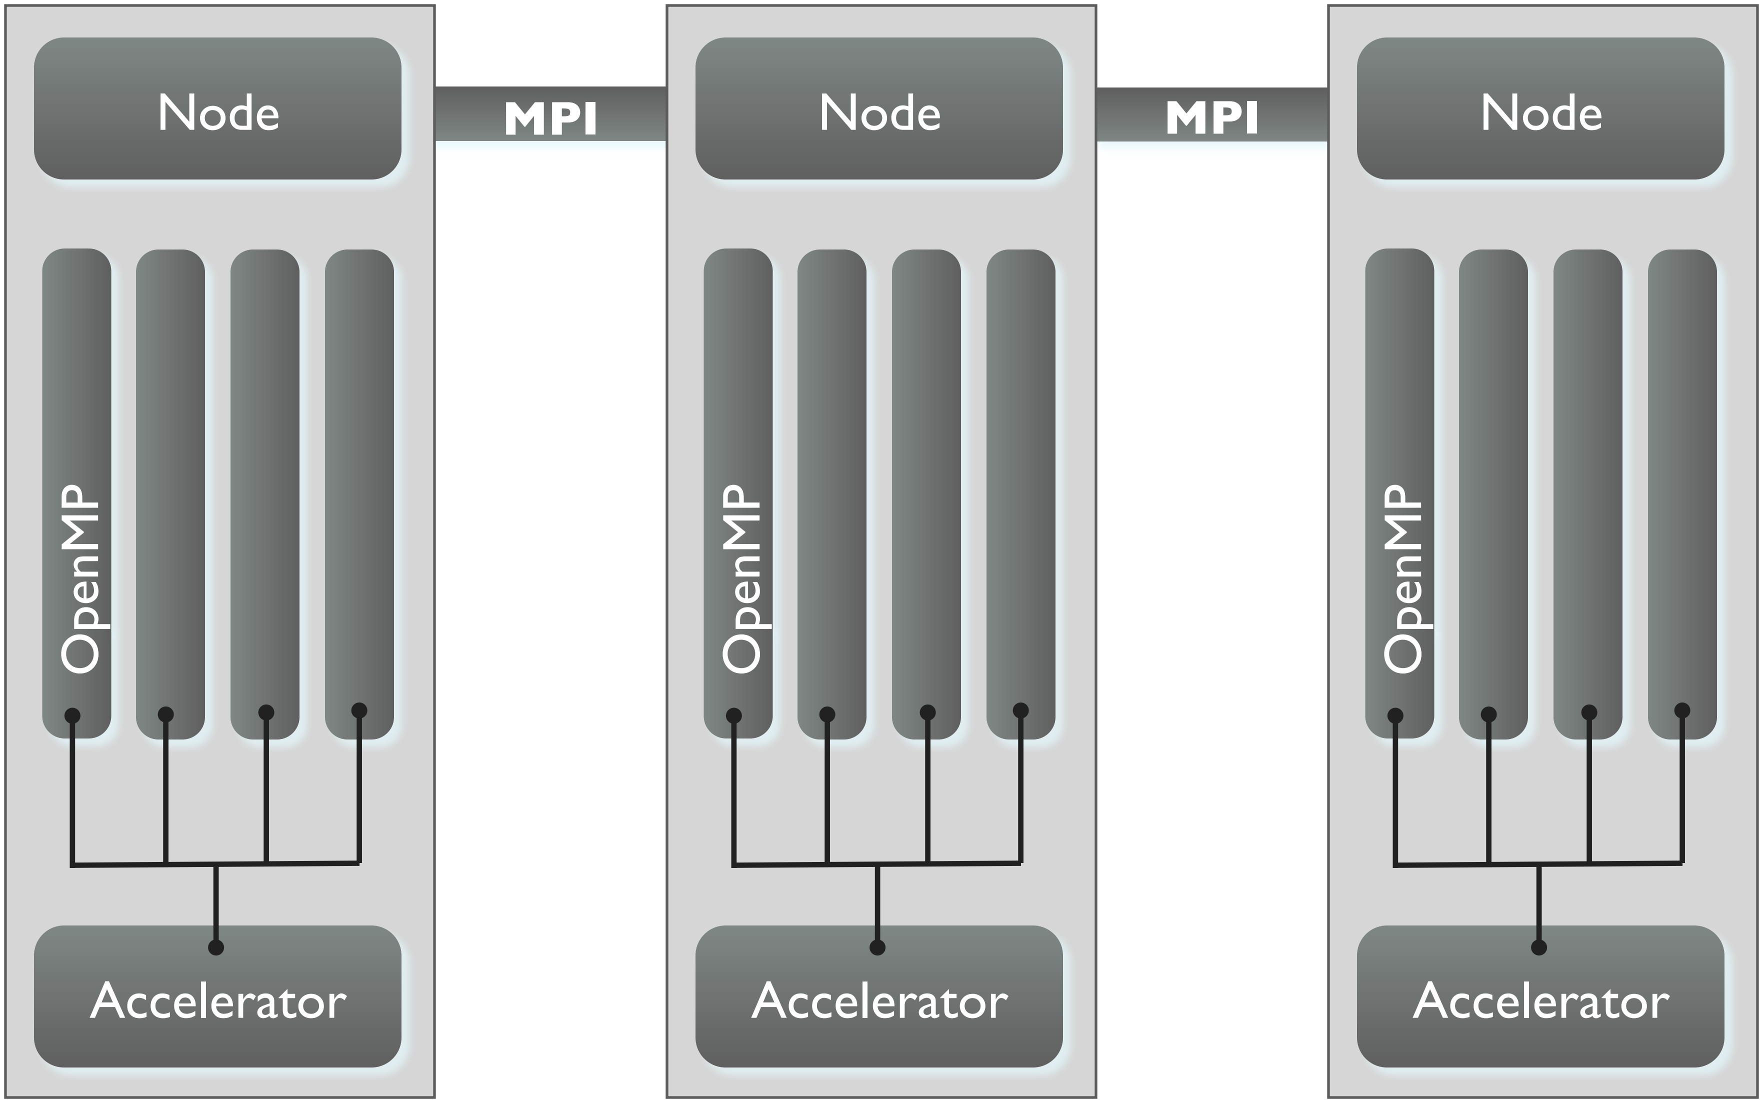

Hybrid Architecture → Scientific Discovery

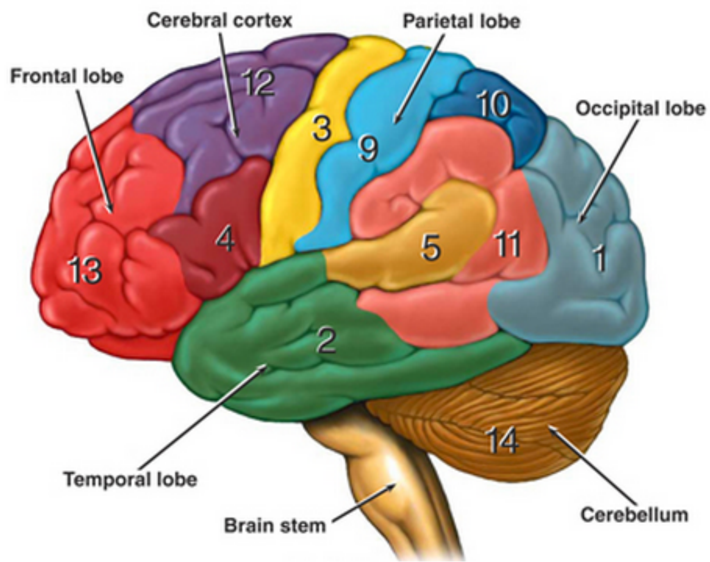

- Evolving hybrid architectures with dedicated centers for specialized tasks surpass solutions from straightforward scaling.
- Significant benefits for those projects that can use it well.
- A growing number of scientific disciplines are benefitting by the resultant speedups.

Early Science Challenges for Titan

WL-LSMS
Illuminating the role of material disorder, statistics, and fluctuations in nanoscale materials and systems.

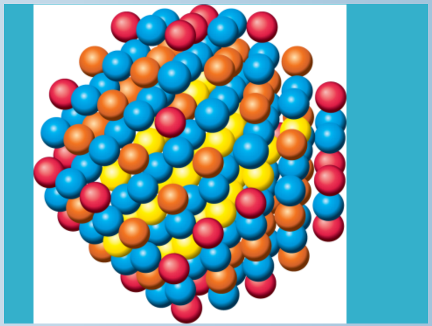

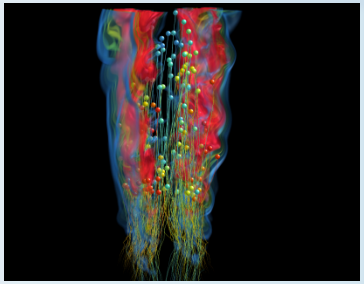

S3D
Understanding turbulent combustion through direct numerical simulation with complex chemistry.
.

NRDF
Radiation transport – important in astrophysics, laser fusion, combustion, atmospheric dynamics, and medical imaging – computed on AMR grids.

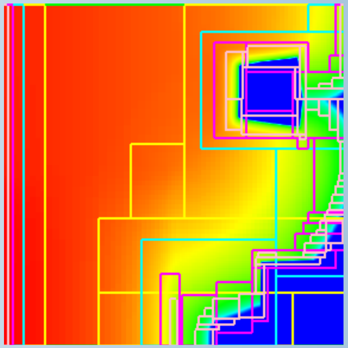

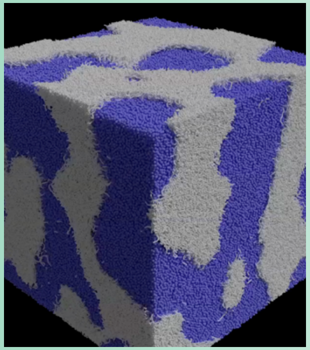

LAMMPS
A molecular dynamics simulation of organic polymers for applications in organic photovoltaic heterojunctions , de-wetting phenomena and biosensor applications

CAM-SE
Answering questions about specific climate change adaptation and mitigation scenarios; realistically represent features like precipitation patterns / statistics and tropical storms.

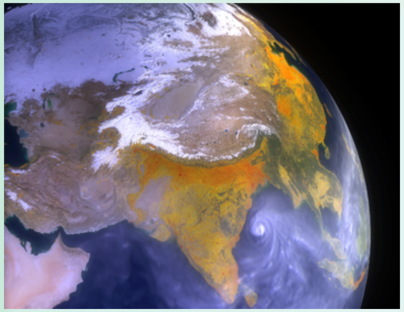

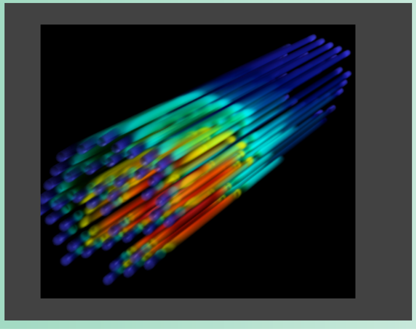

Denovo
Discrete ordinates radiation transport calculations that can be used in a variety of nuclear energy and technology applications.

How Effective are GPUs on Scalable Appications?

OLCF-3 Early Science Codes − *Early* Performance on Titan XK7

Application	Cray XK7 vs. Cray XE6 Performance Ratio*
LAMMPS* Molecular dynamics	**7.4**
S3D Turbulent combustion	**2**
Denovo 3D neutron transport for nuclear reactors	**3.8**
WL-LSMS Statistical mechanics of magnetic materials	**3.5**

Titan: Cray XK7 (Kepler GPU plus AMD 16-core Opteron CPU)
Cray XE6: (2X AMD 16-core Opteron CPUs)
*Performance depends strongly on specific problem size chosen

DOE Computational Facilities Allocation Policy for Leadership Facilities

Primary Objective:

- *"Provide substantial allocations to the open science community through an peered process for a small number of high-impact scientific research projects."*

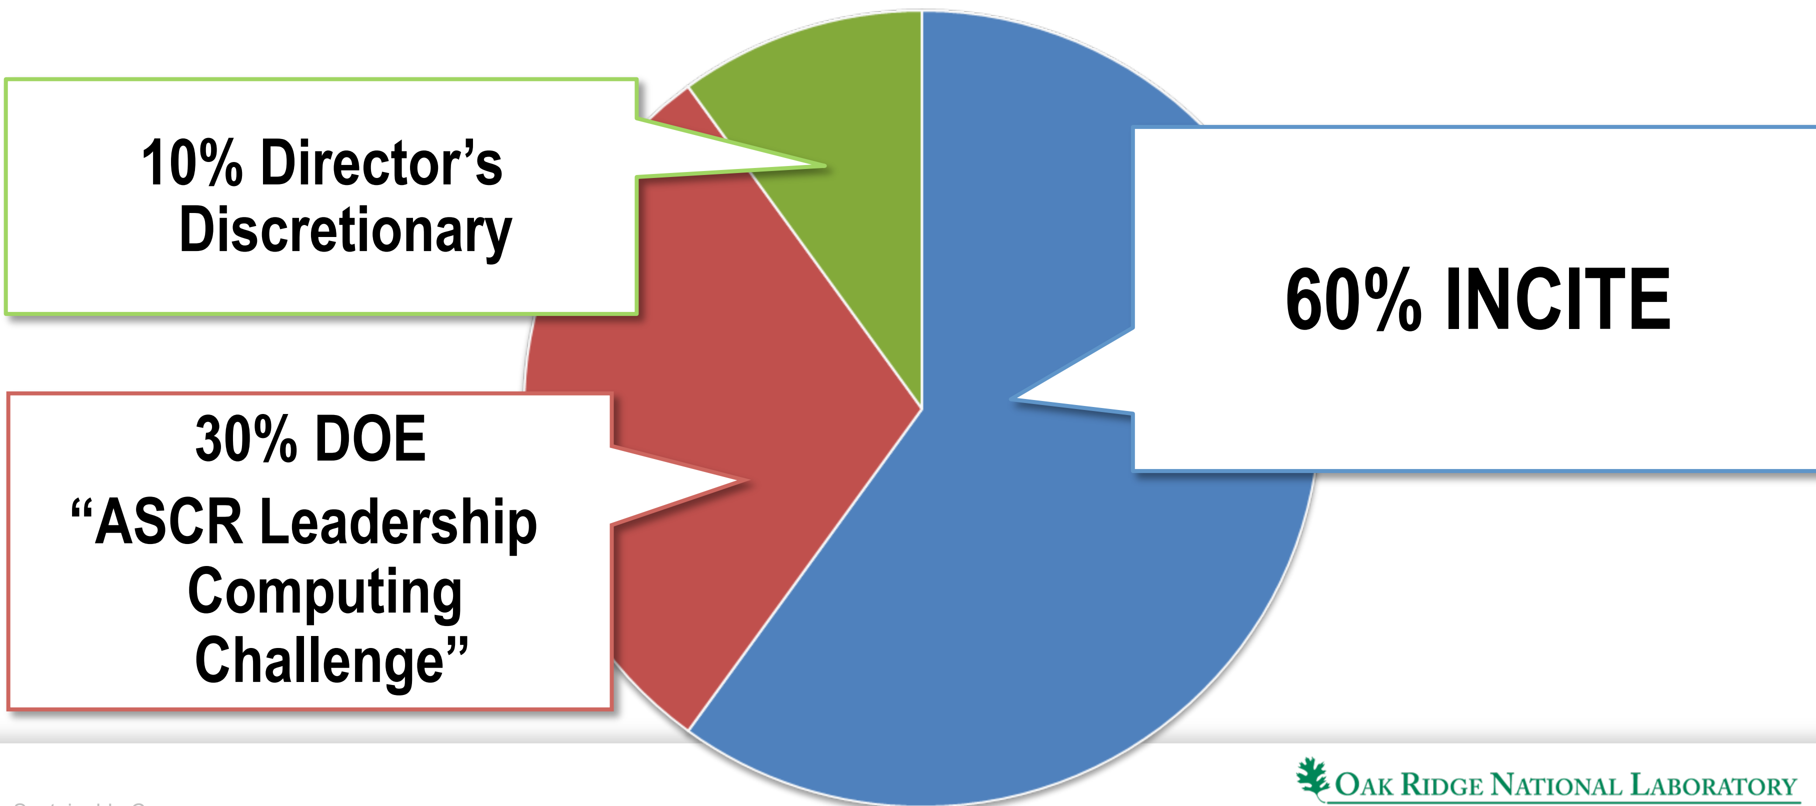

OAK RIDGE NATIONAL LABORATORY
MANAGED BY UT-BATTELLE FOR THE U.S. DEPARTMENT OF ENERGY

Geant4

Transforming Geant4 @ OLCF

20 Years of Excellence in Computational Science

OLCF

Oak Ridge Leadership Computing Facility

- Simulation of the passage of elementary particles through matter constitutes the greatest computational requirement for
  - ALICE and sPHENIX simulations.
  - ATLAS Higgs simulations.
  - nEDM Experiment (at ORNL SNS) simulations.
  - ORNL SNS 2nd Target Station design.
  - Particle beam radiotherapy for cancer treatment.
- See Geant4 talk by R. Mount/ M. Asai.

Computational Jet Tomography

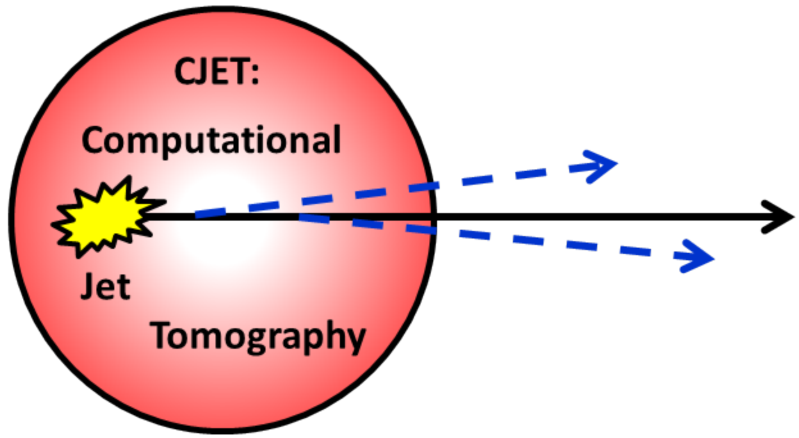

- Understanding quark energy loss in a Quark Gluon Plasma now one of the *primary questions in the field.* Large computations required to simulate predictions from theorist colleagues.
- Project uses new relativistic hydrodynamics code CL-SHASTA, a complete *re-factorization* of CPU-SHASTA using Open CL and best practices of GPU acceleration.
- UrQMD, a leading heavy ion Monte Carlo, uses CPU-SHASTA for transport and will benefit.
- Two publications associated with this project already (PRL acceptance last week).

Open CL SHASTA

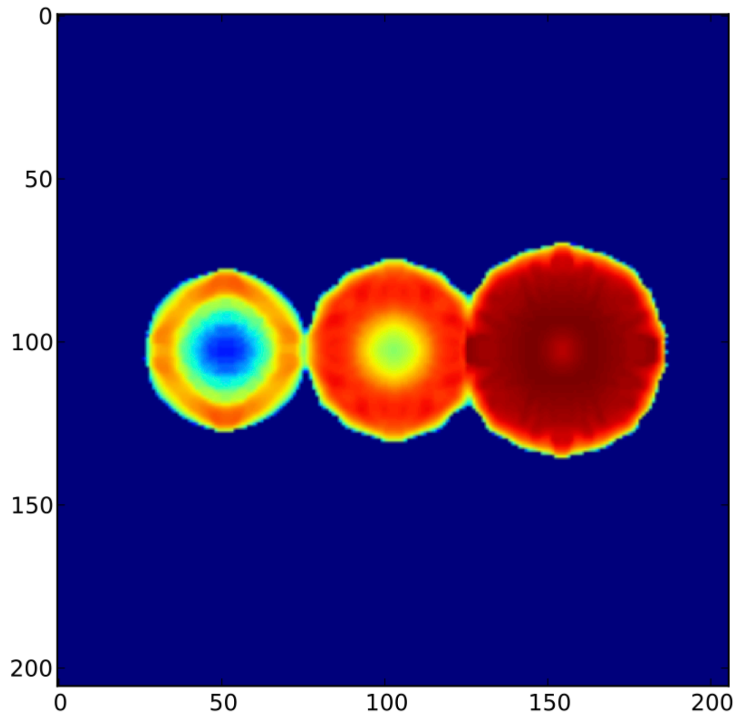

- An exploding fireball from a heavy ion collision with relativistic expansion. CL-SHASTA now running on Titan.
- 11X speedup in port from FORTRAN77 to C++/OpenCL. Subsequent 14X speedup due to GPU acceleration. An overall improvement of 160X.

Open CL SHASTA

ELSEVIER

Computer Physics Communications

Volume 184, Issue 2, February 2013, Pages 311–319

Relativistic hydrodynamics on graphic cards

Jochen Gerhard[a, b], Volker Lindenstruth[a, b], Marcus Bleicher[a, c]

[a] Frankfurt Institute for Advanced Studies, Ruth-Moufang-Straße 1, 60438 Frankfurt am Main, Germany

[b] Institut für Informatik, Johann Wolfgang Goethe-Universität, Robert-Mayer-Straße 11–15, 60054 Frankfurt am Main, Germany

[c] Institut für Theoretische Physik, Johann Wolfgang Goethe-Universität, Max-von-Laue-Straße 1, 60438 Frankfurt am Main, Germany

nEDM at the ORNL SNS

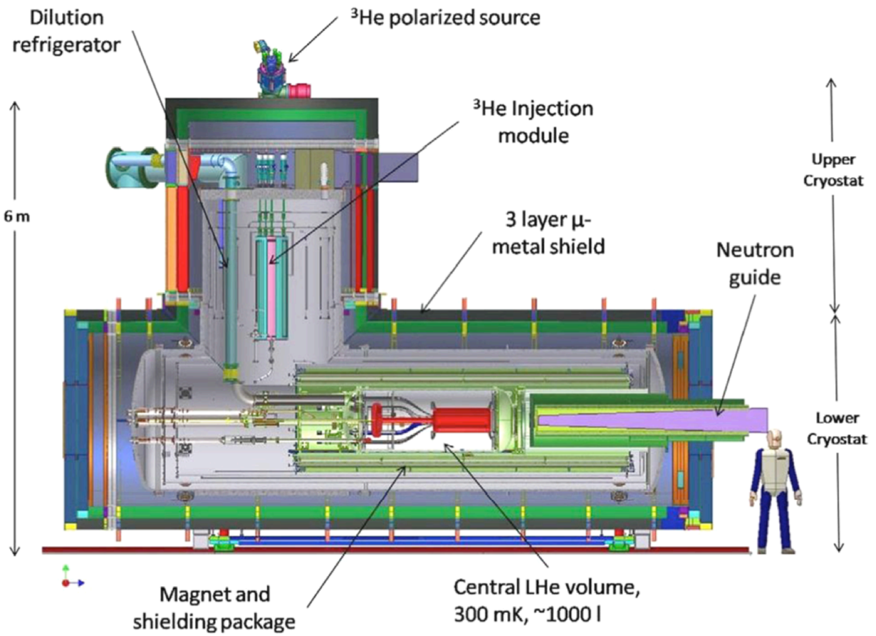

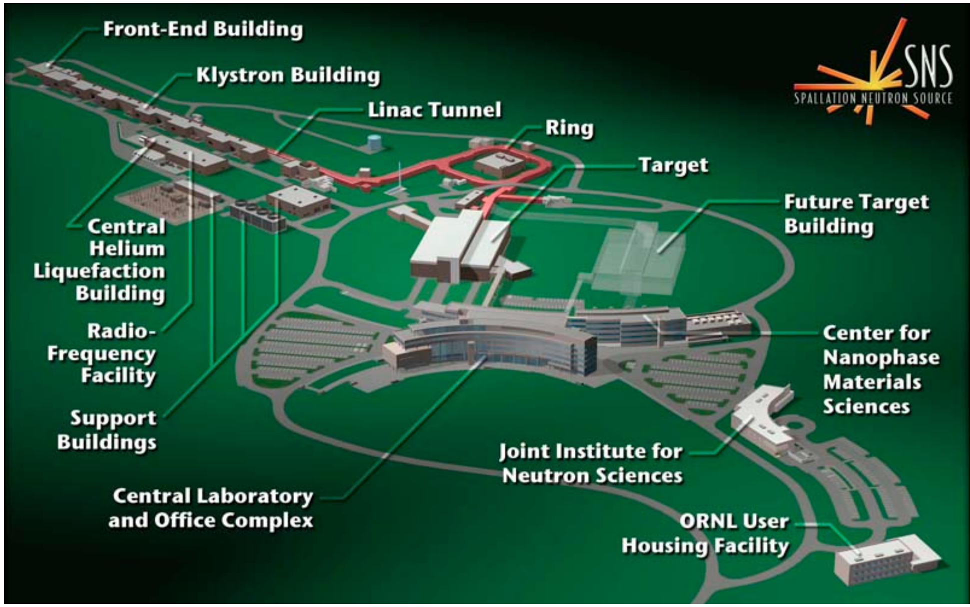

- Neutron Electric Dipole Moment Experiment proposed for fundamental physics beam line at the ORNL Spallation Neutron Source.
- 1M hour Jaguar/ now Titan project in progress to
  - Port and scale nEDM Geant4 simulations on Titan
  - Work with Geant4 experts on low energy neutron transport kernel
  - Facilitate simulation of shielding options for proposed 2nd target station

ORNL SNS

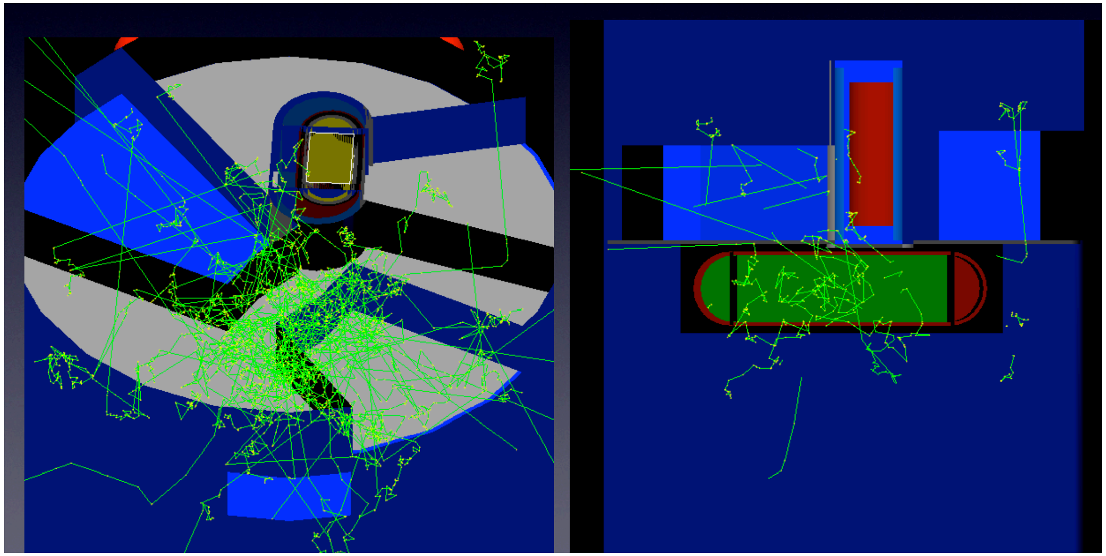

Geant4 simulation of SNS neutron moderators.

- Spallation events simulated. Geant4 is running on Titan *now*.
- Initial G4MPI port achieved. Activation studies performed. Scaling to hundreds of nodes performed.
- Ready to work with Geant4 Steering Board experts for optimization, improved low energy neutron transport, and future acceleration of “low hanging fruit” algorithms likely to benefit from GPU acceleration.

ATLAS PanDA Coming to OLCF

- ATLAS uses sophisticated Production ANd Data Analysis (PanDA) workload management system to optimize data production and availability on the GRID.
- Project underway now to setup and tailor PanDA agent at OLCF.
- This pioneers connection of Titan to the LHC/OSG GRID.
- See talks by S. Panitikin and T. Maeno.

ATLAS PanDA Control *and* Security

- ATLAS needs flexible control, respecting local safety considerations.
- PanDA project reviewed for potential cyber security and export control issues.
- Taking cautious, incremental approach. Trial project approved for initial, limited implementation. Will work to go as far as makes good sense.
- Anticipate significant beneficial use for ATLAS already in 2013.
- ALICE Experiment (without PanDA) very interested in Titan with GRID connectivity and will benefit from this project.

ATLAS is Really Big Data

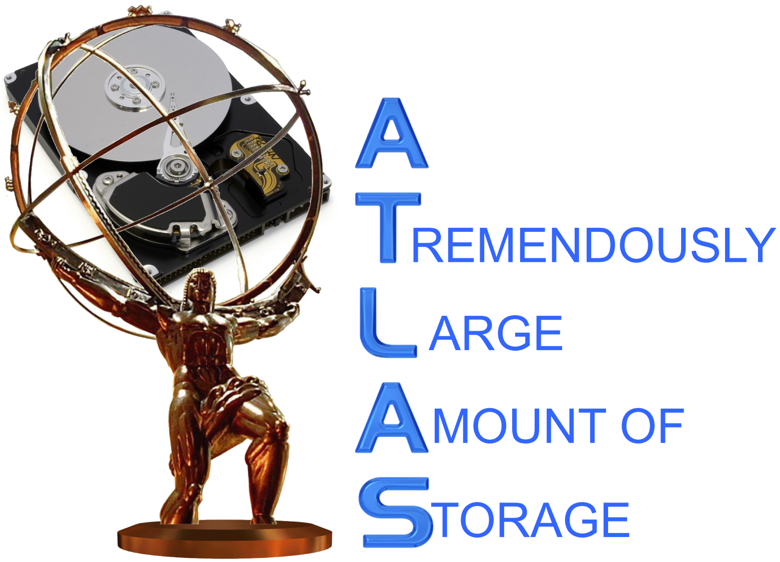

- Solutions for handling Big Data are an OLCF priority.
- OLCF Spider: large, fast Lustre file system. 32 PB and 1 TB/sec. Center-wide, shared resource.
- ATLAS is (one of) the world's *Biggest* Data challenges (recording, storage, retrieval, and transport). 100's of PB and growing.

Got Large Data → August Workshop!

- Great upcoming OLCF training event, August 6 – 8, at University of Tennessee: *Processing and Analysis of Very Large Data Sets*
- On-site or remote registration: https://www.olcf.ornl.gov/training-event/processing-and-analysis-of-very-large-data-sets/
- Would welcome Big Data expert (ATLAS?) to present "lessons learned". Contact foertterfs@ornl.gov (or me) to explore possibility.
- More events here: https://www.olcf.ornl.gov/support/training-events/

Outlook for Particle Physics on Titan

- Today:
  - Geant4 running *now* on Titan. nEDM benefitting from statistics.
  - GPU-accelerated hydro code running.
  - GRID connectivity getting underway in trial mode.
- This year:
  - ATLAS and Geant4 Team benefits and helps pioneer solutions.
  - Titan connected to GRID.
- Over the coming 1 to 2 years:
  - ALICE Experiment and proposed sPHENIX Upgrade will increasingly benefit from these developments.

ROOT on KRAKEN

- KRAKEN is #26 on the Top500 (1.2 PF) and located next to Titan at ORNL. Administered by National Institute for Computational Sciences (NICS), a partner in NSF's Extreme Science and Engineering Discovery Environment (EXSEDE).
- 0.2 M hours exploratory project. Team includes F. Rademakers and P. Hristov (CERN).
- Project goals:
  - Port and scale CERN ROOT software.
  - Port ALICE software.
  - Adoption of the Cray Linux Environment as a supported platform for ROOT and AliROOT.

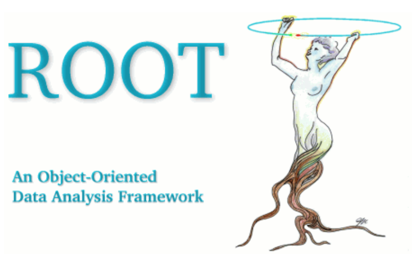

ALICE
A JOURNEY OF DISCOVERY

Conclusions

- Leadership computing is for the critically important problems that need the most powerful compute and data infrastructure. Our compute and data resources have grown 10,000X over the decade, are in high demand, and are effectively used.
- Computer system performance increases through parallelism
  - Clock speed trend flat to slower over coming years
  - Applications must utilize all inherent parallelism
- Accelerated, hybrid-multicore computing solutions are performing very well on real, complex scientific applications.
- OLCF resources are available to academia and industry through open, peer-reviewed allocation mechanisms.
- For further information
  - https://sites.google.com/site/xgeant4
  - https://sites.google.com/site/cjetsite
  - https://sites.google.com/site/openclshasta

Acknowledgements

- Special thanks to J. Wells, OLCF Director of Science, for general OLCF material.
- OLCF-3 Vendor Parters: Cray, AMD, and NVIDIA
- This research used resources of the Oak Ridge Leadership Computing Facility at the Oak Ridge National Laboratory, which is supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC05-00OR22725.